

ARRANGEMENT IN AN IP NODE FOR PRESERVING SECURITY-BASED SEQUENCES BY ORDERING IP PACKETS ACCORDING TO QUALITY OF SERVICE REQUIREMENTS PRIOR TO ENCRYPTION

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

The present invention relates to transport of Internet Protocol (IP) packets, requiring a guaranteed quality of service (QoS), via secure IP connections.

DESCRIPTION OF THE RELATED ART

5 The development of newer protocols for Internet Protocol (IP) networks has extended the capabilities of IP networks. For example, deployment of QoS policies in IP networks has enabled the reliable transport of time-sensitive media data, including audio, video, Voice over IP (VoIP), etc., based on prioritizing the transport of data packets.

10 In particular, data packets identified as associated with latency-sensitive data traffic (e.g., audio, video, VoIP, etc.) are assigned a higher priority than other lower-priority data packets associated with, for example, Simple Mail Transfer Protocol (SMTP) or User Datagram Protocol (UDP) applications. Hence, an outbound interface in a router may include a high-priority queue (for high priority packets) and a low-priority queue (for low priority packets), enabling the packets to be output by the outbound interface according to their priority and determined capacity: if a
15 network interface driver (i.e., an executable software resource configured for controlling the outbound interface) detects backpressure (i.e., congestion), the network interface driver will reorder outbound traffic based on priority of the data packets (e.g., by avoiding outputting packets from the low-priority queue until the high-priority queue has been emptied). Hence, many QoS techniques will reorder packets, causing packets to be output in a sequence that differs from the order the
20 packets were supplied to the outbound interface.

 Problems also exist in maintaining a guaranteed quality of service in cases where a destination host is limited in its downstream capacity. For example, a broadband service provider may limit the downstream bandwidth available to a broadband subscriber; hence, even though a headend router is capable of outputting multiple 1024kbps or higher (e.g., 1.5 Mb/s) media streams on a high-speed

interface (e.g., T1 or higher), the destination router within the broadband network may be configured by the broadband service provider to limit the downstream bandwidth to a contracted rate of 1024 kbps.

Consequently, the destination router may drop downstream packets destined for the broadband subscriber if the amount of data traffic exceeds the contracted rate (e.g., 1024kbps), for example in the case of a destination router receiving, for the broadband subscriber, a 1024kbps media stream plus a burst of SMTP packets. The result is reduced network efficiency due to the unnecessary waste of network resources utilized in generation, routing, and transmission of the data packets that ultimately were dropped by the destination router.

The development of secure IP connections involves IP packets passing through encrypted tunnels. In particular, secure IP tunnels have been used to establish virtual private networks (VPNs) between a local area network (e.g., a corporate LAN) and a remote node (e.g., a telecommuter's computer). In particular, a secure IP tunnel is established between the remote node (referred to as the VPN client) and a VPN server that separates the local area network via the wide area network.

The Internet Engineering Task Force (IETF) has published a Request for Comments (2401), by Kent et al., entitled "Security Architecture for the Internet Protocol" available on the IETF website at <http://www.ietf.org/rfc/rfc2401.txt?number=2401>, the disclosure of which is incorporated in its entirety herein by reference. The above-incorporated RFC 2401 discloses an architecture (referred to as IPSEC) for providing security services for IPv4 or IPv6 data packets at the IP layer., and uses a prescribed Authentication Header (AH) protocol and a prescribed Encapsulating Security Payload (ESP) protocol to provide traffic security. Both the AH protocol and the ESP protocol permit use of anti-replay services (i.e., replay protection), where sequence numbers are added by a transmitting node (e.g., a VPN server) to IP packets being output as a data stream onto an encrypted tunnel.

According to RFC 2401, when a security association (SA) is established between a sender and a receiver, their respective counters (Sequence Counter in the sender and Anti-Replay Window in the receiver) are set to zero: the first packet sent by the sender has a sequence number of "1", the second packet sent by the sender has a sequence number of "2", etc., such that each successive packet output by the sender onto that SA has a corresponding successive sequence number.

Hence, the receiver can expect the received data packets to have a respective contiguous sequence of sequence numbers. If the receiver detects a packet having a sequence number that is out of order relative to a previously received packet, the receiver determines the detected packet is an invalid packet and can discard the packet.

5 The receiver configured for implementing replay protection according to RFC 2401 also will drop packets that are received out-of-order: if the receiver has received packets according to the sequence numbers "1, 2, 3, 4, 100, 101, 5", the receiver will drop the packet having the sequence number "5", since it is out of order relative to the packets having the sequence numbers "100" and "101".

10 As described above, many QoS techniques reorder packets. The IPSEC architecture, in contrast, requires packets to be received in order of the specified sequence numbers to ensure replay protection. Consequently, the inherent inconsistency between QoS techniques and the IPSEC architecture has caused unnecessary packet loss during past attempts at implementing IPSEC protocol and QoS policies on the same router interface.

15 In particular, attempts have been made to add IPSEC functionality to QoS-enabled routers in order to provide latency sensitive traffic (including voice and video) over Virtual Private Networks (VPN). Hence, voice and data packets must pass through encrypted tunnels. To date the voice and data packets have encountered IPSEC encryption and sequence number assignment prior to being passed to the outbound driver that performs the QoS functionality. Hence, any detection
20 of congestion by the outbound driver causes reordering of packets such that the higher priority packets are at the front of the outbound queue.

Consequently, the decrypting peer, having detected an IPSEC sequence number that is out of order, drops the packets that were received out of order, even though the dropped packet is a valid, secure packet.

25 Although some encryption devices utilize queues before input to an encryption chip (i.e., integrated circuit), such queues have been used solely to prevent loss of data due to exceeding the input bandwidth of the encryption chip.

SUMMARY OF THE INVENTION

There is a need for an arrangement that enables latency sensitive traffic to be transported via encrypted tunnels, with guaranteed quality of service, without loss of packets due to reordering of packets bearing sequence numbers.

5 There also is a need for an arrangement that enables transmission of data streams between secure tunnel endpoints with guaranteed of quality service, in a manner that minimizes packet loss between intermediate routers.

10 These and other needs are attained by the present invention, where a router has at least one outbound interface configured for establishing multiple IP-based secure connections (i.e., tunnels) with respective destinations based on transmission of encrypted data packets via the IP-based secure connections. The encrypted data packets are generated by a cryptographic module, where each encrypted packet successively output from the cryptographic module includes a corresponding successively-unique sequence number. The supply of data packets to the cryptographic module is controlled by a queue controller: the queue controller assigns, for each secure connection, a
15 corresponding queuing module configured for outputting a group of data packets associated with the corresponding secure connection according to a corresponding assigned maximum output bandwidth. Each queuing module also is configured for reordering the corresponding group of data packets according to a determined quality of service policy and the corresponding assigned maximum output bandwidth.

20 Hence, the queue controller ensures that data packets are supplied to the cryptographic module in a manner that maintains quality of service policies for latency-sensitive traffic, while ensuring that the data packets are supplied in an order that ensures that quality of services policies implemented by the outbound interface have a minimal likelihood of reordering higher-priority packets. Hence, the queue controller ensures that the data flow remains below outbound interface
25 congestion thresholds, minimizing the need for reordering packets in the outbound interface by QoS based queuing mechanisms. In addition, the queue controller can be configured to ensure that the assigned maximum output bandwidth corresponds to authorized bandwidth rates for a corresponding destination, minimizing the possibility that intermediate routers may drop the packets due to congestion on the downstream link to the destination.

One aspect of the present invention includes a method in a router having at least one outbound interface. The method includes establishing, on the outbound interface, a plurality of Internet Protocol (IP)-based secure connections with respective destinations based on receiving encrypted packets generated by a cryptographic module, each encrypted packet successively output
5 from the cryptographic module having a corresponding successively-unique sequence number. The method also includes controlling supply of data packets to the cryptographic module. The supply of data packets is controlled by: (1) assigning, for each secure connection, a corresponding queuing module, (2) reordering, in each queuing module, a corresponding group of the data packets associated with the corresponding secure connection according to a determined quality of service
10 policy and based on a corresponding assigned maximum output bandwidth for the corresponding queuing module, and (3) outputting to the cryptographic module the group of data packets, from each corresponding queuing module according to the corresponding assigned maximum output bandwidth, for generation of the encrypted packets. The method also includes outputting the encrypted packets from the cryptographic module to the one outbound interface for transport via
15 their associated secure connections.

Another aspect of the present invention includes a router comprising a cryptographic module configured for successively outputting encrypted packets having respective successively-unique sequence numbers. The router also includes an outbound interface, and a queue controller. The outbound interface is configured for establishing a plurality of Internet Protocol (IP)-based secure
20 connections with respective destinations based on receiving respective streams of the encrypted packets. The queue controller is configured for controlling supply of data packets to the cryptographic module, and also is configured for assigning, for each secure connection, a corresponding queuing module. Each queuing module configured for outputting to the cryptographic module a corresponding group of the data packets associated with the corresponding
25 secure connection, and according to a corresponding assigned maximum output bandwidth for the corresponding queuing module, for generation of the corresponding stream of the encrypted packets. Each queuing module also is configured for selectively reordering the corresponding group of the data packets according to a determined quality of service policy and the corresponding assigned maximum output bandwidth.

Additional advantages and novel features of the invention will be set forth in part in the description which follows and in part will become apparent to those skilled in the art upon examination of the following or may be learned by practice of the invention. The advantages of the present invention may be realized and attained by means of instrumentalities and combinations particularly pointed out in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

Reference is made to the attached drawings, wherein elements having the same reference numeral designations represent like elements throughout and wherein:

Figures 1A and 1B are diagrams illustrating routers configured for secure connections based on executing quality of service based queuing, for each secure connection, of packets prior to encryption, according to an embodiment of the present invention.

Figure 2 is a diagram illustrating the queuing controller of Figure 1, according to an embodiment of the present invention.

Figure 3 is a diagram illustrating one of the queuing modules for a corresponding secure connection, according to an embodiment of the present invention.

Figure 4 is a diagram illustrating the method by the queuing controller of Figure 1 of ordering packets for a given secure connection according to quality of service requirements prior to encryption, , according to an embodiment of the present invention.

Figure 5 is a diagram illustrating the ordering of packets by the queue controller of Figures 1A and 1B to minimize anti-replay effects, according to an embodiment of the present invention.

BEST MODE FOR CARRYING OUT THE INVENTION

Figures 1A and 1B are diagrams illustrating routers 12a and 12b, respectively, configured for performing quality of service (QoS) based queueing for each secure connection and prior to encryption, according to an embodiment of the present invention. Figure 1A illustrates a router 12a having multiple outbound interfaces 22, and Figure 1B illustrates a router 12b having a single outbound interface 22 (e.g., a branch router having a LAN-based port and a WAN-based port). Hence, the disclosed embodiment can be implemented in a multiport router 12a or a router 12b

having a single outbound interface 22.

Each of the routers 12a and 12b are configured for establishing secure connections 16 with destination endpoints (i.e., tunnel endpoints) via a wide area network. For example, Figure 1A illustrates a virtual private network (VPN) 10 established between the router 12a and destination endpoints 14, via respective secure connections 16, over a wide area packet switched network 18 such as the Internet. In particular, the router 12a includes multiple outbound interfaces 22 (e.g., OB1, OB2, OB3, etc.), each configured for establishing a plurality of IP-based secure connections 16 (i.e., tunnels) with the respective destination tunnel endpoints 14 via the Internet 18. As illustrated in Figure 1, the outbound interface 22 labeled "OB1" establishes the secure connections 16 labeled "S1", "S2", and "S3" with the destination tunnel endpoints 14 labeled "D1", "D2", and "D3", respectively; the outbound interface 22 labeled "OB2" establishes the secure connections 16 labeled "S4" and "S5" with the destination tunnel endpoints 14 labeled "D4" and "D5", respectively; and the outbound interface 22 labeled "OB3" establishes the secure connections "S6" and "S7" with the destination tunnel endpoints "D6" and "D7", respectively.

The secure connections 16 are established by each outbound interface 22 based on receiving encrypted packets 104 generated by a cryptographic module 20 according to IPSEC protocol. Hence, the destination endpoints 14 could be implemented as a router, a gateway, or a host computer, but in any case the destination endpoints are configured for terminating the secure connection 16 established between the router 12a or 12b and the destination endpoints 14.

In particular, each router 12a and 12b also includes a cryptographic module 20 configured for outputting encrypted packets 104 according to the IPSEC protocol as specified in RFC 2401. Each encrypted packet successively output by the cryptographic module 20 has a corresponding successively-unique sequence number, enabling the destination tunnel endpoints 14 to implement anti-replay procedures according to RFC 2401. As illustrated in Figure 1A, the cryptographic module 20 is configured for outputting: the encrypted packets associated with the secure connections 16 labeled "S1", "S2", and "S3" to the outbound interface 22 labeled "OB1"; the encrypted packets associated with the secure connections 16 labeled "S4" and "S5" to the outbound interface 22 labeled "OB2"; and the encrypted packets associated with the secure connections 16 labeled "S6" and "S7" to the outbound interface 22 labeled "OB3". Figure 1B illustrates that the router 12b is configured

for outputting the encrypted packets 104 associated with the secure connections 16 labeled "S1", "S2", and "S3" to the outbound interface 22 labeled "OB".

Each outbound interface 22 also includes a quality of service (QoS) module 26 configured for implementing a prescribed quality of service procedures in the event that the corresponding
5 outbound interface 22 encounters congestion. In particular, an outbound interface 22 may encounter congestion if the incoming rate of data packets to be transmitted exceeds the available bandwidth on the corresponding outbound link 24. As illustrated in Figures 1A and 1B, the router 12 includes a routing core 34 configured for receiving IP packets from at least one inbound interface 38, and outputting non-encrypted data streams 32. The non-encrypted data streams 32 are illustrated in
10 Figure 1A as "N1", "N2", and "N3" and are output by the routing core 34 to the outbound interfaces 22 labeled "OB1", "OB2", and "OB3", respectively.

If congestion is detected in an outbound interface 22, the corresponding quality of service module 26 is configured to prioritize packets to provide a guaranteed quality of service for latency-sensitive traffic. As described above, however, the prioritizing of packets by the QoS module 26
15 may cause reordering of the encrypted packets output by the cryptographic module 20.

Concerns also arise in the case where a destination tunnel endpoint 14, for example the destination host computer/gateway/router 14 labeled "D1" is limited to a contracted bandwidth rate of 1024kbps: if the destination tunnel endpoint 14 labeled "D1" were to receive downstream video traffic that exceeds the contracted bandwidth rate, then certain downstream packets may be dropped
20 by an access router configured for enforcing the contracted bandwidth rate.

According to the disclosed embodiment, the routers 12a and 12b each include a queue controller 40 configured for controlling supply of data packets (S) that require encryption. In particular, the queue controller 40 is configured for assigning, for each secure connection 16, a corresponding queuing module, described in further detail below with respect to Figures 2-5. As
25 described in further detail below, each queuing module is configured for outputting a group of data packets (e.g., 100a, illustrated in detail in Figures 2 and 5) associated with the corresponding secure connection according to a corresponding assigned maximum output bandwidth. Each queuing module also is configured for selectively reordering the corresponding group of data packets according to a determined quality of service policy and the corresponding assigned maximum output

bandwidth.

Hence, the queue controller 40 ensures that the aggregate output (S') 102 of all the secure connections 14 is less than the input bandwidth of the cryptographic module 20. In addition, each queuing module can be configured to ensure that the packets destined for the corresponding secure connection 16 do not overwhelm the output bandwidth of the corresponding outbound interface 22, and preferably the bandwidth allocated to the subscriber and enforced by the access router 30.

Figure 2 is a diagram illustrating in detail the queue controller 40, according to an embodiment of the present invention. The queue controller 40 includes a security association (SA) assignment module 42 configured for assigning, for each IPSEC-based secure connection 16 (i.e., each Security Association 16), a corresponding queuing module 44. As illustrated in Figure 2, the SA assignment module 42 assigns the queue modules 44a, 44b, 44c, and 44d to the secure connections 16 labeled "S1", "S2", "S3", and "S4", respectively. The flow of packets for the flows associated with the secure connections 16 labeled "S1", "S2", "S3", and "S4" are output by the routing core 34.

The queue controller 40 also assigns a bandwidth controller 46 to each corresponding queuing module 44. Each bandwidth controller 46 is controlling the output bandwidth for the corresponding queuing module (e.g., 44a), hence the bandwidth utilized for the corresponding secure connection 16. Preferably the prescribed threshold utilized by the bandwidth controller 46 is less than the threshold that would be utilized by the quality of service module 26 internal to the outbound interface 22, in order to avoid reordering by the quality of service module 26. Preferably the prescribed threshold utilized by the bandwidth controller 46 also is less than the input bandwidth limit of the cryptographic module 20. Further, each bandwidth controller 46 can be configured such that the sum of bandwidth assigned among the secure connections 16 is less than thresholds for the outbound interface 22 or the IPSEC module 20.

As illustrated in Figure 2, the cryptographic chip 20 performs sequence number assignment (SEQ) to the encrypted packets before being output to the appropriate outbound interface 22.

Figure 3 is a diagram illustrating in further detail an exemplary queue module 44 and the bandwidth controller 46, according to an embodiment of the present invention. Each queue module 44 includes an input controller 50, a packet queue 52, and a connection-specific bandwidth controller

46 configured for outputting to the cryptographic module 20 the group of data packets associated with the corresponding secure connection 16 (e.g., "S1"). In particular, the connection-specific bandwidth controller 46 includes a congestion detector 56, and output controller 58, a maximum output bandwidth register 60, and a bandwidth negotiator 62.

5 The input controller 50 is configured for storing each received data packet that is associated with the corresponding secure connection 16 (e.g., "S1") into one of a plurality of queues (e.g., 52a, 52b, etc.) having respective identified priorities (e.g., "High", "Low", etc.), based on a corresponding identified priority for each packet. As recognized in the art, packets may be prioritized based on packet type (e.g., VoIP, video, TCP, UDP, SMTP, etc.), or some other identifier. Once the packets
10 are stored in the queues 52, the output controller 58 is configured for outputting the stored data packets to the cryptographic module 20 according to the corresponding assigned maximum output bandwidth specified in the register 60. The maximum output bandwidth specified in the register 60 may be manually configured, or may be inserted by the bandwidth negotiation resource 62. For example, the bandwidth negotiation resource 62 may utilize resource reservation protocol (RSVP)
15 in order to communicate with the destination endpoint 14 to identify the maximum downstream bandwidth available to the destination endpoint 14 (e.g., 1024kbps).

 The output controller 58 also can be configured for selectively reordering the data packets stored in the buffers 52a and 52b, relative to the sequence received by the input controller 50, according to a determined quality of service policy in response to a detected congestion condition.
20 The congestion condition may be generated the congestion detector 56. The congestion detector 56 is configured for monitoring the congestion levels in the assigned outbound interface and the levels specified in the output bandwidth register 60. Hence, if any of the above-described congestion levels are detected, the output controller 58 selectively reorders the stored data packets by outputting the data according to the priority queues, such that the high priority queue 52a would be given
25 priority over the data stored in the low priority queue.

 Also note that Figure 3 illustrates that each outbound interface 22 includes an executable driver resource 70 configured for controlling operations of the outbound interface 22, including transfer of data from the cryptographic module 20 and onto the network link 24. Each outbound interface 22 also includes an executable IPSEC resource 72 configured for establishing the secure

connections 16 according to IPSEC protocol, and the QoS module 26.

Figure 4 is a diagram illustrating the method of ordering packets for a given secure connection according to quality of service requirements prior to encryption, according to an embodiment of the present invention. The steps described herein with respect to Figure 4 can be implemented as executable code stored on a computer readable medium (e.g., floppy disk, hard disk, EEPROM, CD-ROM, etc.), or propagated via a computer readable transmission medium (e.g., fiber optic cable, electrically-conductive transmission line medium, wireless electromagnetic medium, etc.).

The method begins in step 80, where the IPSEC resources 72 in each of the outbound interfaces 22 establish the respective secure connections (i.e., Security Associations (SAs)) according to IPSEC protocol. The SA assignment module assigns to each secure connection 16 a corresponding queuing module 44 in step 82. Each queuing module 44 (e.g., 44a) determines in step 84 the corresponding assigned maximum output bandwidth to be used for the corresponding secure connection (e.g., S1). As described above, the assigned maximum output bandwidth may be obtained from a prior manual configuration, or based on the negotiation resource 62 obtaining the corresponding assigned maximum output bandwidth from the corresponding destination 14.

The transmit data (S) that is to be encrypted is routed by the routing core 34 to the queuing controller 40: the SA assignment module 42 forwards each data packet to the assigned queuing module 44 in step 86 based on the secure connection to be traversed by the data packet. The input controller 50 for the queuing module 44 stores each data packet in a selected queue (e.g., 52a, 52b, etc.) based on a determined priority.

The output controller 58 for the queuing module 44 then prepares to output the stored data packets to the cryptographic module 20 in accordance with the maximum permitted output bandwidth specified in the corresponding register 60: if in step 88 the output controller 58 for the corresponding queuing module 44 detects a congestion condition, the controller 58 reorders in step 90 the packets to be output based on priority, for example by outputting from the highest priority queue 52a, minimizing the probability that high priority packets will be dropped. Note that congestion in step 88 can be detected, for example based on the congestion detector 56, or based on the bandwidth controller detecting that levels in the queues 52 approaching congestion levels due

to a higher rate of input data encountered by the input controller 50. The output controller 58 outputs the data packets at or below the assigned maximum output bandwidth to the cryptographic module 20 in step 92.

Note that the selective reordering based on congestion is optional: each queuing module (e.g., 44a) for a given secure connection 16 (e.g., S1) can be configured to reorder all packets according to priority, regardless of the presence of any congestion condition.

Figure 5 is a diagram illustrating the reordering of packets by queuing modules, according to an embodiment of the present invention. The bandwidth controllers 46a, 46b, 46c, and 46d output respective streams of queued packets 100a, 100b, 100c, and 100d, for the respective secure tunnels S1, S2, S3, and S4. For example, the bandwidth controller 46a outputs the stream A1, A2, A3, A4, A5, etc. 100a for the secure connection S1; the bandwidth controller 46b outputs the stream B1, B2, B3, B4, B5, etc. 100b for the secure connection S2; the bandwidth controller 46c outputs the stream C1, C2, C3, C4, C5, etc. 100c for the secure connection S3; and the bandwidth controller 46d outputs the stream D1, D2, D3, D4, D5, etc. 100d for the secure connection S4. The streams 100a, 100b, 100c, and 100d are combined into a combined stream 102 and supplied to the cryptographic module 20 for encryption into encrypted packets 104 (e.g., A1', B1', C1', D1', etc.). The encrypted packets 104 are then supplied to the outbound interface 22: for simplicity, assume that only the single outbound interface 22 of Figure 1B is utilized.

Assume now that congestion encountered by the outbound interface 22 cause reordering of the packets, resulting in the reordered stream 106 output on the outbound link 24. As illustrated in Figure 5, the packets B1' and B2' associated with the stream 100b have been reordered relative to the other packets. However, the prior queuing by the queuing modules 44 on a per-secure tunnel basis ensures that the packets associated with the same secure tunnel remain in the appropriate order. Hence, the encrypted packets (e.g., 100'a, 100'b, 100'c, and 100'd) for a given secure tunnel (e.g., S1, S2, S3, and S4) arrive at the corresponding destination 14 (e.g., D1, D2, D3, and D4) in the appropriate order, minimizing the probability of dropping packets due to anti-replay protection mechanisms in the destination endpoints 14.

According to the disclosed embodiment, reordering of encrypted packets having sequence numbers can be minimized, minimizing the unnecessary loss of processor resources, cryptography

engine resources, and bandwidth resources throughout the virtual private network 10. Further, low priority packets can be dropped or delayed prior to encryption, optimizing resources within the router.

5 While the disclosed embodiment has been described in connection with what is presently considered to be the most practical and preferred embodiment, it is to be understood that the invention is not limited to the disclosed embodiments, but, on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims.